

# Layered Fields for Natural Tessellations on Surfaces

RHALEB ZAYER, Max Planck Institute for Informatics Saarland Informatics Campus, Germany

DANIEL MLAKAR, Graz University of Technology Graz, Austria

MARKUS STEINBERGER, Graz University of Technology Graz, Austria

HANS-PETER SEIDEL, Max Planck Institute for Informatics Saarland Informatics Campus, Germany

Mimicking natural tessellation patterns is a fascinating multi-disciplinary problem. Geometric methods aiming at reproducing such partitions on surface meshes are commonly based on the Voronoi model and its variants, and are often faced with challenging issues such as metric estimation, geometric, topological complications, and most critically, parallelization. In this paper, we introduce an alternate model which may be of value for resolving these issues. We drop the assumption that regions need to be separated by lines. Instead, we regard region boundaries as narrow bands and we model the partition as a set of smooth functions layered over the surface. Given an initial set of seeds or regions, the partition emerges as the solution of a time dependent set of partial differential equations describing concurrently evolving fronts on the surface. Our solution does not require geodesic estimation, elaborate numerical solvers, or complicated bookkeeping data structures. The cost per time-iteration is dominated by the multiplication and addition of two sparse matrices. Extension of our approach in a Lloyd's algorithm fashion can be easily achieved and the extraction of the dual mesh can be conveniently performed in parallel through matrix algebra. As our approach relies mainly on basic linear algebra kernels, it lends itself to efficient implementation on modern graphics hardware.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; *Massively parallel algorithms*;

Additional Key Words and Phrases: Natural tessellations, surface fields, Voronoi diagrams, Lloyd's algorithm

## ACM Reference Format:

Rhaleb Zayer, Daniel Mlakar, Markus Steinberger, and Hans-Peter Seidel. 2018. Layered Fields for Natural Tessellations on Surfaces. *ACM Trans. Graph.* 37, 6, Article 264 (November 2018), 15 pages. <https://doi.org/10.1145/3272127.3275072>

## 1 INTRODUCTION

In his work on celestial bodies, Descartes devised diagrams where planar points (generators, sites) are separated by lines equidistant to their closest neighboring points [Okabe et al. 1992]. The resulting regions enclosed by the lines are called Voronoi cells in honor of Voronoi who studied the problem extension to higher dimensions

Authors' addresses: Rhaleb Zayer, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany, [rzayer@mpi-inf.mpg.de](mailto:rzayer@mpi-inf.mpg.de); Daniel Mlakar, Graz University of Technology, Graz, Austria, [daniel.mlakar@icg.tugraz.at](mailto:daniel.mlakar@icg.tugraz.at); Markus Steinberger, Graz University of Technology, Graz, Austria, [steinberger@icg.tugraz.at](mailto:steinberger@icg.tugraz.at); Hans-Peter Seidel, Max Planck Institute for Informatics, Saarland Informatics Campus, Germany, [hpsidel@mpi-inf.mpg.de](mailto:hpsidel@mpi-inf.mpg.de).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© Association for Computing Machinery.  
0730-0301/2018/11-ART264  
<https://doi.org/10.1145/3272127.3275072>



Fig. 1. Natural tessellations as seen on a tortoise shell (top), drying earth (bottom-left), and a giraffe pattern. Cells are not separated by razor thin lines but rather by narrow bands which are not necessarily straight.

extensively. A restriction which imposes that generators coincide with the center of mass of their respective cells, yields the so called centroidal Voronoi tessellations (CVT). In general, CVT leads to more balanced cells and point distribution as the number of generators increases (Gershon's conjecture [Gershon 1979]). In the planar and volumetric settings, these methods have been used across multiple disciplines ranging from geographic planning to architectural design. Their aesthetic appeal arises from invoking the beauty of natural tessellations, see Figure 1. The direct creation of these tessellations on surfaces however, is inhibited by difficulties arising from estimating metrics on curved surfaces and by the toil of maintaining a valid secondary mesh, geometrically and topologically, which evolves on top of the original one.

In this work, we re-examine the notion of Voronoi diagrams and propose an alternate model which can simplify the creation of natural tessellations on surface meshes. We regard Voronoi diagrams as an idealization of observable natural phenomena. In this idealization, boundaries across regions are modeled as razor sharp lines. In nature, however, this is seldom the case as shown in Figure 1. Furthermore, the natural mechanisms which steer the creation of

these cell-like structures often do not require knowledge of the location of other seeds or the geodesic paths towards them. Different fronts grow simultaneously till they collide and form cell walls. In order to capture this behavior we rely on an idea introduced by Fix to address free boundary problems, and in particular the Stefan type [Fix 1981]. In the numerical modeling of such problems, *e.g.*, melting ice, tracking the location of the sharp interface between the solid and liquid states over time complicates the problem setup. On the other hand, treating the interface as a narrow band where the transition from one state to another is continuous leads to a much more tractable problem formulation.

We drop the assumption that tessellation cells need to be separated by lines. Instead, we regard cell boundaries as narrow bands and we model the partition as a set of smooth functions layered over the surface. Given an initial set of seeds or regions, the partition is obtained as the solution of a time dependent set of partial differential equations describing concurrently evolving fronts. The numerical solution does not require elaborate solvers and is carried out by a simple Euler time-stepping scheme. Furthermore, we show how the problem can be formulated for efficient implementation. In this way, elaborate mesh data structures can be avoided in favor of sparse matrix algebra. Within this formalism, the cost per time-iteration is dominated by the multiplication of two sparse matrices. The extension of our approach in a Lloyd's algorithm fashion [Lloyd 1982] is straightforward. As a proof of concept, using an extremely simple approximation of cell centroids, our approach produces well balanced partitions. Our formulation is intended for producing partitions separated by smooth narrow bands. Nonetheless, sharp boundaries can be obtained at will. Furthermore, our approach allows generating and monitoring the quality of the Delaunay-like dual mesh throughout Lloyd's iterations without compromising performance. As our approach relies mainly on basic linear algebra kernels, it lends itself to efficient implementation on modern graphics hardware and we provide the necessary steps for reproducing our performance.

## 2 RELATED WORK

*Historic developments:* CVT can be achieved by the probabilistic method proposed by MacQueen [MacQueen 1967] which performs generator updates by taking the mean of a randomly sampled point (Monte-carlo method) and its closest generator. A parallel implementation and several improvements of this approach have been proposed by Ju et al. [2002]. Alternatively, an initial Voronoi diagram can be iteratively updated by taking the current configuration centroids as the generators for the next Voronoi diagram [Lloyd 1982]. Reformulation of the problem as an objective function which can be addressed by means of Newton-like approaches has been studied by Iri et al. [1984] as well as Du and Emelianenko [2006].

*Sampling and partitioning of meshes.* The above methods apply to both planar and volumetric settings as Euclidean distance metrics can be directly used and are easy to compute. The extension to surfaces is impeded by the need to measure or approximate metrics on surfaces. This limitation poses a great challenge in computer graphics and steered a steady research effort dedicated to addressing various facets of the problem. In the context of texture generation, Turk [1991] attempts to mimic a reaction-diffusion process on

surfaces and extends it to mesh generation. This idea was further elaborated in [Turk 1992] where a variant of Lloyd's algorithm is developed. It proceeds from a random seed configuration on the surfaces and uses repelling forces to redistribute the seeds evenly. The method requires a local planar unfolding to estimate the repulsion force. By working on implicit surfaces, local flattening can be avoided [Witkin and Heckbert 1994]. The final configuration marks the equilibrium of a set of constrained control points and freely floating particles which can merge or vanish driven by potentials modeled as Gaussian kernels. Our approach, on the other hand, models a growth process driven by variational principles and implicitly and seamlessly handles patch region boundaries.

Although the CVT problem can be addressed by means of restricted Voronoi diagrams, *i.e.*, the intersection of the mesh and a volumetric Voronoi diagram. This approach requires extensive data structure management, may not scale well for large meshes, and faces additional challenges resulting from complicated geometric or topological figures. We do not cover these methods in this brief overview and restrict the discussion to methods which operate directly on surface meshes. In the work of Alliez et al. [2003], the metric problem was avoided by directly flattening the surface and performing Lloyd iterations in the parametric plane. Instead of a global parametrization, the approach of Surazhsky et al. [2003] uses a patch based parametrization to perform Lloyd's updates on the local parametric domain. The approach of Peyré and Cohen [2006] performs CVT on surfaces by taking advantage of the fast marching approach estimation of geodesics [Kimmel and Sethian 1998] which extends the level set approach for solving the eikonal equation [Sethian 1996] to triangle meshes. The authors further use a gradient descent for finding local intrinsic centers of mass. Greedy window reduction schemes such as the VTP approach [Qin et al. 2016] can be used for fast geodesic calculation and Voronoi tessellations [Qin et al. 2017]. However, our experiments with the VTP approach suggest that it is prone to failure on irregular triangulations, *e.g.* Beetle (Figure 8). We presume their window reduction scheme is too aggressive and may not handle special cases correctly.

In a follow-up to Newton like approaches [Du and Emelianenko 2006; Iri et al. 1984], it has been shown that an L-BFGS formulation of CVT is in general faster than the standard Lloyd approach [Liu et al. 2009]. The estimation of CVT on surfaces however is performed by means of standard Euclidean metrics similar to Du et al. [2003] and requires careful face triangle splitting along the boundaries of Voronoi cells. The approach of Xin et al. [2016] extends the latter to power diagrams, and allow generators to evolve on the tangent plane and re-project them on the surface for evaluation of the objective function. The extensive evaluations of geodesics between points [Surazhsky et al. 2005], however restricts the method to very small meshes (1K triangles) and limits its practical scope. Although, performance can be possibly improved by using more recent geodesic approximations [Crane et al. 2013; Ying et al. 2013], it is not expected to drop significantly due to the high number of required queries. The work of Wang et al. [2015] proposes several enhancements to the standard Lloyd algorithm by means of local exponential maps and also extends the same ideas to the framework of Liu et al. [2009].



With a few exceptions, most of existing work avoids direct evaluations on meshes by some sort of global or local planar embedding. The need to compute geodesics, often done by third party code, maintain an evolving mesh topology, and perform several non trivial geometric and topological operations, makes the implementation of these methods a daunting task. The resulting not-so-straightforward algorithmic pipeline complicates the portability of these methods to modern parallel computing hardware. From a simulation standpoint, early methods for producing patterns on surfaces, *e.g.*, Turk [1991] can yield pleasant Voronoi-like patterns, however their lack of controllability and the difficulties in extracting relevant cells inhibits the adoption of these methods as geometric partitioning tools.

Most recently, Herholz et al. [2017] took advantage of the heat diffusion approach [Crane et al. 2013] for approximating geodesic distances to drive the tessellation generation. Although they parallelize their approach on the GPU, they suffer from major scalability issues. Besides the computational and memory costs induced by factorization, they need to determine region affiliation, which requires comparison between the heat diffusion solutions, *e.g.*, for 10 thousand seeds and a mesh of 10 million vertices. Technically this would require cross row comparison in a  $10k \times 10M$  matrix to find the correct approximation of Voronoi Cells. To reduce the cost of comparisons, they estimate a radius for each seed by executing Dijkstra’s algorithm, which by itself has similar complexity to the overall heat diffusion; they ignore the cost of this step in their timings. A similar strategy is adopted by [Qin et al. 2017] which requires quasi-dense sampling provided by other methods. Clearly, such estimates are prone to failure for non-regular seed distributions, *i.e.*, in a linear seed configuration (as *e.g.* shown in Figure 4) limiting the diffusion distance to the closest  $n$  other points is not sufficient, as the final cells do not limit each other but rather stretch out over the entire domain. In order to obtain region boundaries, additional postprocessing which includes subdivision and topological cleanups is required. In contrast, our approach grows all cells concurrently and thus, stops the growth exactly and only when a cell’s growth is blocked along its entire boundary.

### 3 LAYERED FIELDS

As a motivating example, consider the problem of describing the state  $\varphi$  of a melting ice cube. The shape of the ice is continuously evolving over time. Modeling the interface across the liquid and solid states as a sharp boundary gives rise to several numerical challenges pertaining to discontinuities and to tracking the exact location of the boundary at each time step. Instead, when the boundary between the two states is defined as a narrow band where the function  $\varphi$  describing the evolution gradually varies, these problems can be avoided and the location of the boundary is simply inferred from the state function [Fix 1981; Steinbach 2000]. In our context, the evolution of a growing region, or in other words, the position of its boundary, as a function of time, is implicitly given by the evolution of its state function  $\varphi$  which is set to 1 inside the region, 0 outside and it is continuously varying as it traverses the narrow boundary band separating them as illustrated in Figure 2. An important advantage of this formalism is that there is no need to track the boundary, *i.e.*, to explicitly describe its location via mathematical equations, during




Fig. 2. Definition of the state function over a region, and its profile as it crosses through the boundary (grey) of the region (green).

evolution. In this way, issues pertaining to sharp boundaries, such as derivative continuity, are avoided in the first place.

#### 3.1 Derivation

We regard the problem of tessellation formation as a growth process initialized at a set of given seeds (sites), or initial regions, and driven by the evolution of the growing regions boundaries (interface). In view of Figure 2, the interface can be described as a set of evolving narrow bands whose motion is consistent with the mechanism that governs the overall growth process. It is therefore possible to describe the entire tessellation by a single function, and to extract the interface after equilibrium as the values confined strictly between 0 and 1. Our exploration of this direction revealed that extracting the individual cells after convergence requires additional intricate geometric and topological operations. Instead, we describe the state of each cell by associating a function  $\varphi_i$  with each seed  $i$ . In this way, we can capture the state of the entire tessellation while having easy access to the individual cells as illustrated in Figure 3.

More formally, assume we have a set of  $n$  evolving regions, initialized by seeds, over a given surface. We associate a function  $\varphi_i$  with each region. We define  $\varphi_i$  as a function which takes value 1 inside the cell and 0 outside the cell. At the front, we have  $0 < \varphi_i < 1$ , see Figure 2 for an illustration of the case of a planar surface.

In the following we establish some rules which govern the interaction between the fields  $\varphi_i$  across the different layers. From the natural consideration that at equilibrium, any given point on the surface can only be inside a single region, or the narrow band separating different regions, we impose a partition of unity condition at all times

$$\sum_{i=1}^n \varphi_i = 1. \quad (1)$$

Clearly, before equilibrium is reached, there will be surface areas that are neither inside a growing region or on a narrow band, for this reason we introduce a base layer which is initialized to 1 for all points. When seeds are initialized, the corresponding locations in the base layer are set to 0.

In order to drive the evolution of our model, we need to define its energy. Although the regions evolve on different layers, *cf.* Figure 3, their interiors should not spatially overlap to reflect that cells are not allowed to merge into a single cell. This can be regarded as the fields associated with different cells being mutually exclusive. It ensues that we need to account for the products  $\varphi_i \varphi_j$  for equilibrium which implicitly guaranties that each field is localized within a finite compact region.

Since the state functions of type  $\varphi_i$  are constant within a region, they have null gradients and do not drive the evolution of the region.




Fig. 4. The progress of the propagation from seeds on two different planar configurations with periodic boundary conditions. Note that the effect of the  $g$  as shown in Figure 5 is visualized in blue.




Fig. 3. Illustration of the concept of layered fields. Each field evolves on a different layer, in this way, cell information is easily accessible by querying the corresponding layer.

In contrast, the gradients on the narrow boundary bands do, *cf.* profile in Figure 2. To reflect the condition that a growing region should not be allowed to change the shape of an already existing region, we can, similarly to above reasoning, capture the kinetics of the cell boundaries through the dot products of their gradients. In other words, we need to account for  $\nabla\varphi_i\nabla\varphi_j$  for equilibrium which allows distinguishing fields in the boundary regions.

It remains to describe the behavior of the system when two boundary bands meet. Let the corresponding energy contribution be denoted as  $g$  for now. We will define it shortly.

From the discussion above, we define the global energy of the model on the surface  $S$  as a combination of three terms,

$$F = \int_S f ds = \int_S \sum_{i=1}^n \sum_{j=i+1}^n (w_{ij}\varphi_i\varphi_j - \frac{a_{ij}}{2}\nabla\varphi_i\nabla\varphi_j + g) ds, \quad (2)$$

where the penalty  $w_{ij}$  and gradient energy  $a_{ij}$  terms are fixed scalars. In all our experiments (we keep a ratio of 5, with  $w_{ij} = .2$ ). We foresee them as parameters that could be used to adapt our approach to natural processes such as biological cell growth.

By introducing the Lagrange multiplier  $\lambda$  to account for the constraint of equation 1, the  $\varphi_i$ 's can be treated as independent variables,

and the Lagrangian then reads

$$\Gamma = \int_S f + \lambda \left( \sum_{i=1}^n \varphi_i - 1 \right) ds. \quad (3)$$

The progression of  $\varphi_i$  in the direction of the minimum of  $\Gamma$  can then be expressed as follows

$$\dot{\varphi}_i = -\frac{\partial\Gamma}{\partial\varphi_i} = -\frac{\partial F}{\partial\varphi_i} - \lambda. \quad (4)$$

In the above notation, we consider only the direction and we drop the time scale (for now). In order to neutralize the Lagrange multiplier, consider the following intermediate function

$$\phi_{ij} = \varphi_i - \varphi_j; \quad (i < j). \quad (5)$$

Clearly,  $\phi_{ij} = -\phi_{ji}$ . Substituting into equation 1, we obtain

$$\varphi_i = \frac{1}{n} \left( \sum_{j=1}^n \phi_{ij} + 1 \right). \quad (6)$$

As the time evolution of  $\phi_{ij}$  can be written in virtue of equation 4 as

$$\dot{\phi}_{ij} = \dot{\varphi}_i - \dot{\varphi}_j = -\frac{\partial F}{\partial\varphi_i} + \frac{\partial F}{\partial\varphi_j}. \quad (7)$$

it follows that the time evolution of  $\phi_{ij}$  is independent of the Lagrange multiplier  $\lambda$ . From equation 6, we obtain

$$\dot{\varphi}_i = -\frac{1}{n} \sum_{j=1}^n \left( \frac{\partial F}{\partial\varphi_i} - \frac{\partial F}{\partial\varphi_j} \right). \quad (8)$$

By considering a time scale (*i.e.*, switching to standard time derivatives), equation 8 can be written as

$$\dot{\varphi} = -\frac{1}{n} \sum_{j=1}^n \mu \left( \frac{\partial F}{\partial\varphi_i} - \frac{\partial F}{\partial\varphi_j} \right), \quad (9)$$

by means of constant mobility term scalar  $\mu$ .

On the other hand, the functional derivative  $\partial F/\partial\varphi_i$  reads

$$\begin{aligned}\partial F/\partial\varphi_i &= \frac{\partial f}{\partial\varphi_i} - \nabla \cdot \frac{\partial f}{\partial\nabla\varphi_i} \\ &= \sum_{\substack{j=1 \\ j\neq i}}^n (w_{ij}\varphi_j + a_{ij}^2\nabla^2\varphi_j) + \frac{\partial g}{\partial\varphi_i}.\end{aligned}\quad (10)$$

Please note that in the above expression, the vector in the denominator (gradient in this case) is a standard shorthand notation for directional derivatives. By means of equations 9 and 10, we obtain

$$\begin{aligned}\dot{\varphi}_i &= -\sum_{j=1}^n \frac{\mu}{n} \left( \sum_{k=1}^n \left[ (w_{ik} - w_{jk})\varphi_k + \frac{1}{2}(a_{ik}^2 - a_{jk}^2)\nabla^2\varphi_k \right] \right. \\ &\quad \left. + \left( \frac{\partial g}{\partial\varphi_i} - \frac{\partial g}{\partial\varphi_j} \right) \right).\end{aligned}\quad (11)$$

From the expression above, it turns out we do not need to define  $g$  explicitly. Instead, we can directly define the expression of the difference term  $(\frac{\partial g}{\partial\varphi_i} - \frac{\partial g}{\partial\varphi_j})$ . To motivate our choice, we recall that  $g$  encodes the desired behavior when two boundary bands




Fig. 5. Plot of  $\sqrt{\varphi_i(1-\varphi_i)}$ .

meet each other. In the simple case where only two bands  $\varphi_i, \varphi_j$  meet, we have by virtue of equation 1,  $\varphi_j = 1 - \varphi_i$ . In practice, we seek a function that produces a symmetric behavior in the interval  $(0, 1)$  and vanishes at 0 and 1. One such function is  $\sqrt{\varphi_i(1-\varphi_i)}$  which has a suitable profile as illustrated in Figure 5.

In this way, we define the remaining term

$$\frac{\partial g}{\partial\varphi_i} - \frac{\partial g}{\partial\varphi_j} = -\sqrt{\varphi_i\varphi_j}e_{ij}\quad (12)$$

where  $e_{ij}$  is a constant scalar which describes the strength of interaction between the two boundary bands.

This completes the definition of the time derivative of the governing equation of the system. The solution can then be carried by a simple explicit Euler stepping scheme

$$\varphi_i(t + \Delta t) = \varphi_i + \dot{\varphi}_i\Delta t.\quad (13)$$

To provide a first intuition of our approach, we present initial experiments for two different seed configuration in the planar setting in Figure 4. The experiments are in accordance with what would be obtained using Euclidian metric for partitioning and confirms the well behaved nature of our simple approach. The effect of not accounting for gradient interactions at the narrow bands is shown in Figure 6. Extending our approach to the spatial setting, we focus on the particular case of the sphere where a Delaunay triangulation for a set of points laying on the sphere can be gained through computing their convex hull [Renka 1997]. Starting from the same initial seed configuration shown as the green dots in Figure 7), the partitioning obtained with our approach is quasi-similar to the one from convex hull based approach and confirms the meaningfulness of our approach. For visualization, we use the squared sum of the field values at every vertex.




Fig. 6. Effect of omitting the  $g$  term in Equation 11, would remove the separation between the fields.




Fig. 7. Side by side comparison of our approach (left) with the convex hull based Delaunay triangulation approach to Voronoi diagrams on the sphere [Renka 1997] (right). The original seeds are shown in green. A zoom-in reveals the quasi-similarity of the results.

### 3.2 Efficient problem representation

In this section, we devise a concise and efficient representation of the problem geared towards simple implementation and portability. We represent the fields as a sparse matrix of size  $(n_r \times n_v)$ , where  $n_r$  is the number of seeds or partitions and  $n_v$  the number of vertices in the whole mesh. For each vertex  $v_i$ , we associate the corresponding values of  $\varphi$ s across the layers, i.e.,  $\varphi_{v_i} = (\varphi_{1v_i}, \varphi_{2v_i}, \dots, \varphi_{n_r v_i})$ . Since a vertex will either belong to a single region, will lie on the front across two different regions, or even at a junction (triple, quadruple, ...),  $\varphi_{v_i}$  is mostly populated by zeros. With this remark in mind, we represent the whole multi-layered field as the sparse matrix

$$\Phi = (\varphi_{v_1}, \varphi_{v_2}, \dots, \varphi_{v_{n_v}}).\quad (14)$$



From a data querying perspective, this formalism offers an elegant way of extracting topological information. Vertices within a cell are the ones that flag non-zero values along the row in  $\Phi$  that corresponds to the cell.

From an implementation perspective, this formalism simplifies the algorithmic pipeline to a great extent. Since the geometry of the mesh is static, we pre-compute the Laplacian  $L$  of the mesh only once. At each iteration, the evaluation of the second term of equation 11 on the whole surface and over all layers amounts to a single sparse matrix-matrix multiplication, *i.e.*,  $\Phi L^T$ .

The algorithmic outline of our approach is shown in Algorithm 1 and can be broken down into four major steps:

- (1) initialization based on seeds or regions (ln 1)
- (2) Laplacian computation and evaluation of the layered field Laplacian (ln 2-4)
- (3) update of the layered field according to equation 11 (ln 7-14)
- (4) normalization (ln 20)

---

**ALGORITHM 1:** Multi-layered field algorithm
 

---

```

1 Initialize  $\Phi$  with seeds and/or initial regions
2 Precompute Laplacian  $L$ 
3 for  $t \leftarrow t_0$  to  $t_{end}$  do
4    $L_t = \Phi L^T$ 
5   for  $i \leftarrow 1$  to  $n_v$  do
6      $irows = (\Phi_{*,i} > 0) \vee ((\Phi_{*,i} = 0) \wedge (L_{t*,i} > 0))$ 
7      $n_i = |irows|$ 
8     for  $j \leftarrow 1$  to  $n_i$  do
9        $d \leftarrow 0$ 
10      for  $k \leftarrow 1$  to  $n_i$  do
11         $sum \leftarrow 0$ 
12        for  $l \leftarrow 1$  to  $n_i$  do
13           $sum += \frac{1}{2}(a_{j,l} - a_{k,l})L_t(l, i) + (w_{j,l} - w_{k,l})\Phi(l, i)$ 
14           $d += -\frac{\mu}{n_i} \left( sum - e_{j,k} \sqrt{\Phi(j, i)\Phi(k, i)} \right)$ 
15       $\Phi(j, i) += d \Delta t$ 
16      if  $(\Phi(j, i) > 1)$  then
17         $\Phi(j, i) = 1$ 
18      if  $(\Phi(j, i) \leq 0)$  then
19         $\Phi(j, i) = 0$ 
20 Normalize each column of  $\Phi$  by the sum of its nonzeros

```

---

*Initialization.* Our approach can be initialized by setting the values of the field  $\Phi$  at the corresponding vertices to 1. Ideally, a vertex and its immediate neighbors can be set to 1 or to a function defined on its neighboring triangles. There is no need at this stage to accommodate the boundary of the seeded regions to account for the narrow band of the interface. Please note, that energy defined earlier accounts mainly for interaction between cells. In order to allow cells to evolve freely beforehand, we introduce an additional base layer as mentioned earlier in Section 3. The corresponding  $\varphi_b$  evaluates to 1 at all non-seeded vertices.

*Laplacian computation.* Since the mesh is static, we precompute the Laplacian of the mesh and store the resulting sparse matrix. Since the Laplacian is not updated throughout the rest of our algorithm, the cost of this operation does not affect the overall iterative process. Nonetheless, within the time loop the Laplacian of the phase field needs to be updated (ln.4.), the cost of this operations amounts to a sparse matrix-matrix multiplication.

*Update.* The main operation within the time loop is to update the layered field  $\Phi$ . This spans lines 5 – 20. The first step is to detect the growing regions, that is regions where the field is nonzero as well as possible areas that can get affected by the propagation, *i.e.*, regions where the field Laplacian is positive. Both operations are performed in ln. 6. The output of this operation is  $irows$  which indicates the active rows for the field evaluated at vertex  $i$ . For the columns  $\varphi_i$ , we iterate over the rows of interest, *i.e.*,  $irows$  and perform a summation over the entries of the corresponding field Laplacian multiplied by a factor which encodes the nature and the strength of interactions across the involved layers. Clearly, this operation is akin to sparse matrix-matrix addition with additional scaling for the involved quantities. Once this columns addition is completed, a cleaning operation needs to be performed to keep the field within the defined range, namely, above zero and below one.

*Normalization.* Once the field update is computed, a normalization is performed to enforce that the sum of field values across layers at every vertex equals 1.

From the discussion above, the cost of one time step is dominated by two operations and can be regarded as the cost of sparse matrix-matrix multiplication and a sparse matrix-matrix addition. The latter is in fact relatively negligible. Given the pure algebraic nature of our approach it can be easily ported to modern parallel computing platforms. An outline of our implementation on the graphics computing unit (GPU) is given in the following section.

## 4 GPU IMPLEMENTATION

The major challenge of a GPU implementation is finding a good way to parallelize the entire algorithm. The first step in each iterations is a simple sparse matrix-matrix multiplication, for which a variety of GPU libraries exist. Mostly iterating over the matrix columns, we use a compressed sparse column (CSC) representation, which describes a matrix using a linear array of values and associated row indices, as well as a column pointer array that points to the beginning of every column. The second step, the update, consists of four nested loops (ln 5, 8, 10, 12). The two outer loops describe updates of individual entries of a sparse matrix (ln 15). As these operations are independent of one another, these loops form a major potential for parallelization. From a high-level perspective this corresponds to a sparse matrix-matrix addition. However, this update is complicated by the inner loops (ln 10 and 12), which each require serial iterations over entire matrix columns. Also, the parallelization of the outer two loops is non trivial, as the loops run over the entries of interest, which vary between vertices and also change from iteration to iteration. The final step of each iteration, the normalization, can obviously be parallelized over each column.

*Laplacian.* Although sparse matrix-matrix multiplication is algorithmically the most complicated step on a parallel device like the GPU, we can make use of readily available libraries like cuSparse [NVIDIA 2016]. Internally, cuSparse computes an explicit transpose for  $L^T$ , which can take up nearly half the overall time spent on the multiplication. As we consider the Laplacian of the mesh constant throughout all iterations, we can precompute the transpose once, and reuse it throughout all iterations.

The Laplacian of the phase field  $L_t$  is computed in every iteration. Allocating memory for it in every iteration would lead to a significant overhead. Thus, we allocate memory before starting the iterations and reuse it in every step. However,  $L_t$  can obviously grow from iteration to iteration. Thus, in case the preallocated storage is not sufficient for  $L_t$ , we have to allocate it anew. To provide enough storage for multiple iterations, we enlarge it by 20% at once.

*Update.* In order to parallelize the for loops involved in the update routine, we first determine the values of  $i$  and  $j$  throughout the entire update. To this end, we need to identify all rows of interest (irows) for all vertices. We do this by evaluating Ln 6 of the algorithm in parallel for all vertices (columns of the involved matrices). The result of this operation corresponds to a vector irows for every vertex, which as a whole corresponds to a sparse matrix skeleton—describing where rows of interest are, without assigning any value to them. We call this matrix skeleton the *interest skeleton*. The construction of the skeleton involves a prefix sum [Sengupta et al. 2007] over the number of entries in every column. Memory needs to be allocated for the skeleton in every iteration. We again, follow the same strategy as before, allocating memory once and enlarging it by 20% in case it is not sufficient.

After the interest skeleton has been constructed, we transfer  $\Phi$  and  $L_t$  into a representation that matches the skeleton, adding explicit zero entries for entries not present in the original matrices. We call these representation  $\hat{\Phi}$  and  $\hat{L}_t$ . Although mathematically there is no difference between  $\Phi$  and  $\hat{\Phi}$ ,  $L_t$  and  $\hat{L}_t$ , having all involved matrices within the same skeleton (and thus the same non-zero pattern), allows for efficient parallelization of the ensuing steps.

Our next concern is the innermost summation of the algorithm (ln 12-13). Analyzing its structure, it can be seen that it describes an interaction between the different fields at a given vertex. When evaluating the influence of field  $k$  on  $l$  and field  $l$  on  $k$  equal terms are considered. A straight forward implementation would compute these respective continuations multiple times. Taking a slightly different approach, we can avoid these duplicated computations. The loop in line 12 with the partial sum contributions can be reinterpreted as the differences of selected results from the matrix-matrix multiplications

$$A_L = a \cdot L_t \quad \text{and} \quad W_\Phi = w \cdot \Phi.$$

Using this result, the summation (loop in ln 12) can be rewritten:

$$sum = \frac{1}{2}(A_L(j, i) - A_L(k, i)) + (W_\Phi(j, i) - W_\Phi(k, i)), \quad (15)$$

whereas only selected results from the respective matrices are used. As these used values correspond exactly to the entries present in the interest skeleton, we can compute them efficiently in parallel. Starting one thread for every entry in the skeleton on the GPU, we

assign each thread to the row  $i$  querying its own row coordinate from the skeleton. We then compute a dot product between a dense vector (row  $i$  in  $a/w$ ) and a sparse vector (the entire column associated with the current entry, representing  $(l)$ ) to compute the resulting entry of  $A_L/W_\Phi$ .

With  $A_L$  and  $W_\Phi$  at hand, we compute the update for  $\Phi$  (ln 14) for all entries of  $\hat{\Phi}$  in parallel. Again starting one thread for each entry in the interest skeleton, each thread runs through the loop over  $k$  (ln 10). We evaluate the sum according to equation 15 by fetching four values from  $A_L$  and  $W_\Phi$  and compute the update to  $d$  (ln 14) by fetching the respective values from  $t$ ,  $e$  and  $\hat{\Phi}$ . Note that due to using  $\hat{\Phi}$  here, the lookup for  $\Phi(j, i)$  and  $\Phi(k, i)$  is trivial although the matrices are sparse, because all involved matrices follow the interest skeleton. Furthermore, the individual entries  $\Phi(k, i)$  correspond to a single column in the interest skeleton and thus are right next to each other in memory.

The update for  $\Phi$  and clamping of the value (line 15 – 19) can be completed alongside the computation of the update itself. The output of this entire step is an updated copy of  $\hat{\Phi}$ .

*Normalization.* The normalization is carried out via two simple steps. First, we compute a sum over each column in parallel. Second, we launch one thread for each individual entry of  $\hat{\Phi}$  and update each value. Working with  $\hat{\Phi}$  instead of  $\Phi$  may introduce additional zero entries in the result (note that ln 18 and 19 may also introduce explicit zero entries). It is desirable to remove those entries to avoid unnecessary increases in the representation. However, we do not have to perform this compaction step explicitly, as it is implicitly carried out by computing the interest skeleton in the next iteration.

After computing the final iteration, we perform an explicit compaction to avoid zero entries in the output. We combine this compaction with the normalization step, counting the number of non zero entries in every column alongside its sum. We then again run a prefix sum to generate the compacted column pointer array as required by the sparse matrix representation and fill the compacted matrix while multiplying with the normalization factors.

*Typical relative performance.* To provide an intuition about the costs of the described steps, we provide detailed timings for one iteration of the algorithm on the Bimba model depicted in Figure 17 with 3.7M triangles and 5000 seeds in Table 1. Clearly, the performance is dominated by the matrix-matrix multiplication. After computing the interest skeleton, the remaining steps can be parallelized very efficiently, leading to a negligible cost for the update itself.

Matrix-Matrix Multiplication	20.241ms	84.3%
Interest Skeleton Creation	1.321ms	5.5%
Computing $\hat{L}_t \hat{\Phi}$	0.720ms	3.0%
Updating $\hat{\Phi}$	1.177ms	4.9%
Normalization	0.552ms	2.3%
<b>Sum</b>	<b>24.011ms</b>	<b>100%</b>

Table 1. Performance breakdown for a single iteration of our algorithm on the Bimba model (Figure 17) with 3.7M triangles and 5000 seeds for an NVIDIA GTX Titan Xp GPU. Due to our optimizations on all other steps, the matrix-matrix multiplication clearly dominates performance.

## 5 LLOYD’S ALGORITHM AND MESH EXTRACTION

Adapting our approach to perform Lloyd-like iterations is straightforward. Starting from a random distribution of sites on the surfaces, we perform the first partitioning as outlined earlier. The vertices of each individual cell are simply the ones that flag nonzero entries in the row of the sparse matrix  $\Phi$ , corresponding to the cell layer. For performing center of mass update, one can resort to existing methods discussed in Section 2. Interestingly though, we observed that even a very simple approximation of the centroid location works well with our approach. We compute the center of mass as the average of the barycenters of all triangles of the cell (have a vertex in the cell) weighted by their respective triangle area. Clearly, this center of mass does not necessarily lay on the surface. We back-project it on the cell along the vector obtained as the average of the triangle normals weighted by their respective area. Cells shapes can be complicated, and at times, the intersection can be more than a single point. In this case, we simply choose the closest to the center of mass. If there is a ray miss, we do not update the seed position at that particular iteration. In this way, cell-center computation does not require explicit distance computation and the overall procedure adheres to the simplicity sought in this work.






Fig. 8. Lloyd-like iterations on the Beetle mesh ( $2M\Delta$ ) for 1000 (left) and 5000 (right) seed locations. The initial random distribution (first row) features extremely uneven clustering of seeds (right). The distribution becomes significantly better after a first approximate center of mass update (second row) and after 15 Lloyd-like iterations (last row).

While more elaborate schemes for center of mass estimation could be used, *e.g.*, [Peyré and Cohen 2006], this simple procedure works well in all our experiments, as shown in Figure 8, and serves the



(a) Bimba 5000



(b) Flightsuit 10 000

Fig. 9. Initial cell area histograms and the effect of 1 and 10 Lloyd iterations for the models shown in Figure 17 and Figure 14.

purpose of illustrating the feasibility of Lloyd’s iteration with our layered-fields representation. The evolution of cell area distribution summarized in Figure 9 confirms the well behaved nature of our approach.

### 5.1 Cell neighborhoods and dual mesh

There are many ways for determining cell neighborhoods and extracting the dual mesh. For instance, tracing along a given levelset of the function defined as the squared sum of  $\Phi$  across the layers would be an option. In fact, the resulting isolines inside the narrow band can be used in a similar fashion to field-based remeshing methods, *e.g.*, [Schall et al. 2008]. However, we found that our layered field representation offers a more malleable and efficient alternative.

We define isolines per individual cells. In view of the definition in equation 12, a threshold value .25 or slightly lower for the field  $\Phi$  would amply guarantee intersection of neighboring cell isolines and subsequently that their underlying triangulations overlap. An example of such isolines is shown in Figure 10a as blue closed curves (a much lower threshold would simply yield too much overlap). However, instead of performing an extensive isoline computation and intersection search over the whole mesh, we capitalize on the overlap itself to define the cell-cell adjacency relations.

In our formulation, the layered field is represented as a sparse matrix where each column represents a vertex and each row represents a cell. A preliminary estimation of the cell-cell adjacency matrix can be obtained by assuming that two cells are adjacent if they share a vertex. Such matrix can be obtained as  $A_v = \Phi\Phi^T$ , with



the diagonal set to zero.  $\bar{\Phi}$  is the sparse matrix where values of  $\Phi$  that pass the threshold are set to one and to zero otherwise. Unfortunately  $A_v$  cannot capture all cases induced by the intersections of cell-cell isolines as two overlapping regions can intersect without sharing a vertex, as shown in the zoom in Figure 10a. Those cases need to be captured at the triangle level. To do so, we can obtain a different adjacency which encodes that cells are adjacent if their respective isolines pass on the same triangle. It can be obtained through sparse matrix algebra as well using a logical variant of  $A_t = (\bar{\Phi}\bar{M})(\bar{\Phi}\bar{M})^T$ , where  $\bar{M}$  is the  $n_v \times n_f$  face-vertex incidence matrix with  $n_v$  and  $n_f$  being the number of vertices and faces respectively. This is the sparse binary form of the recently introduced mesh matrix formulation [Zayer et al. 2017]. Again for  $A_t$  we set the diagonal to zero.

Clearly, the nonzero entries of  $A_v$  are a subset of those of  $A_t$ . Furthermore, entries which are unique to  $A_t$  encode cases where two or more isolines travel the same triangle. However, they do not reveal whether they effectively intersect. We can single out these cells using the nonzero entries of  $(A_t - A_v)$  and obtain the candidate triangles on which the isolines will be effectively computed and the intersection checks performed. These triangles are obtained based on overlapping entries of the respective rows of the involved cells in the matrix  $(\bar{\Phi}\bar{M})$ . In practice, these constitute only a very small number of triangles. For instance, there are only two such cases (red circles) on the mesh in Figure 10a. The zoom-in shows two sets of triangles (black and red) and their corresponding isolines (two shades of green and blues resp.). Once this check is performed, we can confirm whether the involved isolines intersect at those triangles by simple intersection checks. In this way, we can update the corresponding entries of  $A_t$  according to the intersection scores. If there is no effective intersection the entry is set to zero. This adjacency matrix is not guaranteed to yield a manifold mesh as there could be more than three isolines that intersect at certain locations, as shown in example Figure 10b. As the treatment of such degeneracies is generally difficult in post processing, we overcome this issue by automatically enforcing the requirement on those candidates in the adjacency matrix itself. For the candidates in  $(A_t - A_v)$ —when an intersection is detected—we check if the two involved cells have two or more already confirmed shared neighbors. If so, we drop the additional intersection in  $A_t$ . Throughout this simple approach complicated topological cleaning is avoided upfront.

## 5.2 Mesh construction

After the manifoldness check above, we can proceed to construct the dual mesh from the adjacency matrix  $A_t$ . To the best of our knowledge, existing graph based approaches require a planar embedding and then perform ring construction based on angle considerations. Our approach is simpler and is based on matrix algebra as described below. Starting from the curated version  $A_t$ , for a given vertex  $i$ , the corresponding vertex ring (neighbors) are the nonzero elements of row  $i$  of  $A_t$ . Popping one element from the ring, say vertex  $j$ , we simply intersect the neighbors of  $j$  with the ring elements. This can be also algebraically encoded as the multiplication of the rows of  $A_t$  corresponding to the ring elements with the vector  $x_i$  which is set to zero everywhere and one at  $i$ . This step yields a vector with a




Fig. 10. (a) Segmented cells and their lax isolines with a threshold of .25 for the Tweety model. Some cell intersection only happen on triangles but not on vertices (red), with the zoom-in detailing one such intersection. (b) Even after exact isoline intersection, non-manifold meshes would be generated if more than three isolines intersect. (c) We prohibit this situation from happening by restricting the number of shared neighbors of pairs of potentially intersecting cells.

maximum of two nonzeros entries, each of which forms a triangle with  $i$  and  $j$ . We repeat this process till the ring is empty. Once all vertices have been processed an initial triangulation is obtained.

A typical case of spurious triangles may occur when a cell has only three neighbors which are also neighbors to each other (around a bump or a spike). These case can be easily discarded by finding rows of  $A_d$  which contain only three entries and checking if these three entries form a triangle. This last step can be conveniently done by multiplying those rows by the binary mesh matrix of the newly constructed dual mesh. Rows in the output vector which flag a value of 3 refer to the ids of spurious triangles which need to be removed from the triangulation. The last step in the dual mesh construction amounts to consistently orienting the triangles. This can be done in the customary way starting from an initial triangle and orienting its edge-neighbors consistently and repeating the process till all triangle have been processed.

## 6 CONVERGENCE AND STABILITY ANALYSIS

In our current numerical formulation we adopted a basic Euler step approach. Besides simplicity, this choice is motivated by two main reasons: suitability for the problem at hand and performance gains. Due to its simplicity, it enables our approach to achieve high performance for large meshes sizes—which to the best of our knowledge have not been reached by existing approaches.

The use of an implicit stepping approach such as backward Euler or the Adam-Bashforth approach would require use of iterative or direct solver within each step. This clearly affects performance and limits the size of problems that can be addressed especially if a direct solver is used due to fill-in effect in the factorization. Iterative solvers on the other hand may require a large number of iterations and/or adequate preconditioning. Furthermore, the extension to centroidal Voronoi diagrams where dynamic updates are performed would not be as straightforward as with our approach. In this respect, we would like to emphasize that the approach proposed by Herholz et al. [2017] is not an implicit variant of ours. In fact, it is not governed by a time dependent equation and their solution is not an implicit approach at all since they solve for each cell separately. Additionally, their cells




Fig. 11. (top) The Beetle model ( $2M\Delta$ ) and a zoom-in highlighting irregular variation in density and mesh sizing. (bottom) Sharp visualization (field with the largest influence) of the convergence rates for different time steps 0.1, 1.0, 5.0, and 10.0 for 1200, 100, 40, and 30 iterations, respectively. Even for this highly irregular mesh, there is hardly any difference after convergence for  $\Delta t = 0.1 - 5.0$ , only at  $\Delta t = 10$  differences are noticeable. Please zoom in for details in the electronic version.

are not the result of a free evolution but rather of an initial estimate of prescribed radii using Dijkstra's shortest path algorithm.

While a theoretical stability analysis for the governing equations of our current system (equations 11, 12, 13) is beyond the scope of this work given the difficulty of such exercise for nonlinear problems, see for instance LeVeque [2002]. Our empirical tests suggest that the problem is not stiff and for varying time steps the resulting tessellations are remarkably identical. A very challenging case is summarized in Figure 11 for a the Beetle model which features a highly irregular triangulation with varying densities. As can be




Fig. 12. The initial tessellation (left) obtained from an increasing number of seeds (500, 1000 from top) randomly distributed over the Sforza model ( $500K\Delta$ ) and the resulting tessellation after 100 Lloyd-like iterations.

seen, a variation of the time step across a factor of 50 achieves very similar results. Only, when increasing the time step further to  $\Delta t = 10.0$  visible differences show up when extracting the exact border between cells. Note that the number of iterations required does not significantly reduce when going from  $\Delta t = 5.0$  to  $\Delta t = 10.0$ , which indicates, that  $\Delta t$  is getting closer to the usable maximum, as triangles always require one iteration to propagate the field. Nevertheless, the results are still remarkably similar.

## 7 RESULTS

Throughout our experiments, we used the following hardware configuration: an Intel i7 6800k CPU with 32GB of memory and an NVIDIA Geforce Titan Xp with 3840 compute cores and 12GB of memory. We used  $\Delta t = 5.0$ .

*Field propagation and Lloyd-like iterations.* A performance breakdown of our approach and an overview of the used models is given Table 2. Renderings of the models are found in Figure 8, 12, 17, 16, and 14. Detailed step timings are given in Figure 13. As can be seen, the per-step timings increase with each iteration until the fields converge. This fact can be attributed to  $\Phi$  growing in size as the fields propagate, which increases the cost of the sparse matrix-matrix multiplication. As expected, both, the mesh size as well as the number of seeds influence the step timings. However, as the number of seeds on a mesh increases, fewer steps are required until the fields converge (also indicated by the graph for higher seeds counts ending early in Figure 13). Thus, the overall time until field convergence can even reduce with increasing number of seeds.

The spikes in the steps timings indicate reallocation of either of the matrices and thus also an increase in memory consumption. Nevertheless, our approach can handle very large meshes, like the



Model	Faces	Seeds	$t_{min}$	$t_{max}$	$t_{mean}$	$t_{conv}$	Memory
Tweety	54k	500	1.22	6.95	2.44	366	49
		2000	6.21	45.03	19.99	296	184
Sforza	500k	1000	2.45	12.59	4.56	672	197
		5000	2.57	30.87	12.52	1250	443
Hand	1.5M	1000	5.05	20.27	8.46	2544	247
		5000	5.67	37.81	15.20	2251	705
Beetle	2.0M	1000	6.21	20.39	10.40	3132	361
		5000	6.88	38.10	18.13	2718	940
Bimba	3.7M	5000	10.73	51.15	24.01	4793	1091
		7500	11.76	55.35	27.29	4122	1689
Orchid	4.0M	5000	12.06	58.24	25.76	5129	1513
		7500	12.94	60.27	29.39	4443	2031
Flightsuit	21.4M	7500	62.30	139.55	106.92	4761	3542
		10000	64.89	157.52	125.13	3064	4838

Table 2. Performance comparison for various test models and seed numbers. Timings are provided in *ms*;  $t_{min}$ ,  $t_{max}$ , and  $t_{mean}$  are the respective step timings for one Euler step;  $t_{conv}$  is the average time (*ms*) for one Lloyd-like iteration, *i.e.*, the sum of step timings until the fields are converged. Memory lists the peak memory consumption during the executed steps in Mb.




Fig. 13. Detailed step timings in *ms* for the first five Lloyd iterations with different numbers of seeds (500, 1000, 2500, 5000, 7500) for the Sforza and Bimba models.



Fig. 14. Earhart's Flight Suit ( $10M\Delta$ ) with 10 000 seeds, the resulting cells are visualized in different colors. A close-up of the cell distribution (top-row) and the underlying mesh (middle and bottom rows) reveals small scale geometric complexity (pocket fold, button fold) that can successfully be processed with our computationally- and memory efficient approach.

Amelia Earhart Flight Suit depicted in Figure 14. The maximum memory consumption is directly related to the mesh size and the number of seeds.

Typical examples of our experimental results are shown in Figures 8, 12, and 17. The images show the results for an initial random distribution and the results after a given number of Lloyd-like iterations. Our approach works on highly detailed models as the Amelia Earhart Flight Suit, as well as very complicated geometric figures such as the Eembreea Orchid model (Figure 16). The Bimba model (Figure 17) is interesting because it features holes and mesh irregularities in the ear, mouth, and hair regions. Nevertheless our results remain well behaved. Our use of sparse matrices for encoding the layered formulation of our approach offers the flexibility of accessing the individual regions and extracting smooth and sharp boundaries as illustrated in Figure 11 and 14.

The accompanying video shows that our approach can handle complicated geometric figures and large models without any additional requirements. All parameters defined in our work were kept unchanged throughout all experiments. As our fields propagates directly on the surface, we do not require additional snapping or re-projection to remain on the surface as commonly used in other approaches.

The examples in Figure 15-top feature cases with low seed count. Our result for the sphere with 6 seeds reflects a well balanced partition obtained from an initial random distribution. The result for the






Fig. 15. Typical results on topologically simple models (top row) with low seed count and complex models (lower rows) with increasingly higher seed count. The initial random seeds configuration is shown for the sphere (6 seeds), the Tweety (25 seeds), Fertility (500 seeds), Hand (1K seeds), Hip Bone (5K seeds) and Dancing Children (10K seeds) and their corresponding equilibrium throughout Lloyd's iterations.

Tweety model substantiates that complex geometric features can be handled correctly even when the patch size is relatively large. The additional results on more geometrically and topologically complex models confirm the robustness of our simple centroid update procedure to patch size variation.

In Figure 18, we show side by side comparison of our approach to the diffusion diagram approach proposed in [Herholz et al. 2017] throughout multiple iterations. We can readily observe that our approach preserves features (see mouth, hair braids) more faithfully. In addition, the topological features holes and mesh irregularities near the ear, mouth, and hair regions are handled correctly and do not produce any holes or sampling problems. It can be also observed that the method of [Herholz et al. 2017] tends to smoothen out features as it inherits the behavior of an iterative Laplacian process. For the Orchid model, we can see that features tend to appear and disappear in the results of [Herholz et al. 2017], for instance, the most left petal and Orchid stigma, whereas the tessellation behavior




Fig. 16. The progress of the propagation from an initial random distribution of 7500 seeds on the Eembreea Orchid model ( $4M\Delta$ ) and the resulting tessellations after one and 100 Lloyd-like iterations and the extracted dual mesh with  $15k\Delta$ .

is much more consistent in our approach. These observations are also confirmed by the quality metrics in the accompanying table.

*Dual-mesh extraction.* Dual-mesh extraction examples are shown in Figure 16 and 17. Table 3 shows quality metrics computed for the different test cases and various seed constellations. Successful Lloyd iterations ideally redistribute the seeds in such a way that the extracted dual meshes match the original model more closely. At the same time, the mesh quality should also increase. As can be seen in the figures and in the tables that is the case. The mean Hausdorff distance between the extracted dual-mesh and its variance decrease with increasing iterations. Furthermore, the triangle quality (based on the ratio between circumcircle and incircle radius) increases. This goes hand in hand with the minimum angle increasing, the maximum angle decreasing and the number of angles below  $30^\circ$  lowering. We recorded consistent behavior throughout all tested meshes.

*Large seed count.* Use cases with a large count of seeds are shown in Figure 19. All models feature highly detailed geometry and/or interesting topological properties. Our approach handles these cases correctly as can be observed by visual inspection and efficiently as shown in the memory and performance data. Please note that for large models, we simply assign the constant values of the parameters  $a_{ij}$ ,  $e_{ij}$ ,  $w_{ij}$  by a function instead of storing them as a large matrix, in this way, the memory usage is reduced significantly.

## 8 CONCLUSION

The overall objective of this work was to show that by changing the way we apprehend natural tessellations on surfaces a different numerical model, which breaks away from the centuries old straight line diagrams, is possible. In this model, tessellations emerge as the solution of a system of time dependent partial differential equations layered on the surface.

Lloyd iterations	Tweety 100			Sforza 1000			Hand 1000			Beetle 5000			Bimba 5000			Orchid 7500		
	0	1	21	0	1	61	0	1	79	0	1	81	0	1	77	0	1	93
Dist (%)	1.14	0.90	0.89	0.20	0.15	0.14	0.38	0.21	0.83	0.10	0.08	0.07	0.09	0.05	0.04	0.11	0.06	0.06
Dist RMS (%)	1.42	1.19	1.19	0.35	0.26	0.23	0.63	0.37	0.26	0.19	0.14	0.11	0.15	0.09	0.08	0.20	0.11	0.11
Mean Quality	0.62	0.87	0.91	0.80	0.86	0.93	0.82	0.87	0.96	0.77	0.84	0.93	0.81	0.87	0.94	0.64	0.79	0.89
Min Quality	0.11	0.37	0.43	0.01	0.13	0.41	0.03	0.20	0.63	0.01	0.31	0.62	0.02	0.06	0.34	0.01	0.05	0.26
Mean Min Angle	30.62	38.87	45.85	30.16	39.89	47.16	30.24	39.69	53.45	26.87	36.61	45.64	39.24	43.17	51.11	29.16	35.38	45.32
Min Angle	12.3	19.8	37.0	3.8	14.5	37.4	5.3	18.5	44.1	1.3	8.6	28.8	0.9	8.5	29.2	0.4	2.4	24.9
< 30° (%)	44.80	4.59	0.0	52.97	5.68	0.0	54.70	4.61	0.00	60.16	12.10	1.17	19.69	5.49	0.11	56.24	24.00	5.06

Table 3. Quality metrics for the dual meshes extracted from various models for increasing number of Lloyd-like iterations (till convergence). Dist corresponds to the Hausdorff distance between the original and the dual mesh relative to the model’s diagonal (lower is better); Quality is the average ratio between circumcircle and incircle radius for all triangles (1.0 is best); Mean Min Angle corresponds to the average over the minimum angle in each triangle (60° is best).




Fig. 17. Five selected iterations of our Lloyd-like algorithm on the Bimba model (3.7MΔ) and the extracted meshes. Clearly, the extracted meshes increase in quality and get closer to the original model as the number of iterations increases.

The emphasis of this work was on supplying the underlying numerical machinery which allows for representing and manipulating the resulting cells efficiently. As a proof of concept, we showed throughout a simple extension to Lloyd’s-like iterations that intricate geometric and topological operations pertaining to cell updates and dual mesh extraction can be translated into simple linear algebra formulations. Despite this simplicity, our solution is comprehensive and has minimal requirements.

Channeling the numerical effort to linear algebra kernels allowed for taking advantage of fine grained parallelism of modern GPUs. In this way, models encompassing millions of faces can be processed at unprecedented rates allowing for direct monitoring of various aspects of the tessellation process.


## ACKNOWLEDGMENTS

Images in Figure 1 are courtesy of Laszlo Ilyes (Shell Games), Neil Palmer (cracked earth), and Rouven Kurz (Giraffen). All under the creative commons license. Mesh data sets are courtesy of the Smithsonian Institution, Stanford University, and Aim@shape Repository.

## REFERENCES

- Pierre Alliez, Éric Colin de Verdière, Olivier Devillers, and Martin Isenburg. 2003. Isotropic Surface Remeshing. In *Proceedings of SMI '03*. 49–58.
- Keenan Crane, Clarisse Weischedel, and Max Wardetzky. 2013. Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow. *ACM Trans. Graph.* 32, 5, Article 152 (Oct. 2013), 11 pages.
- Qiang Du and Maria Emelianenko. 2006. Acceleration schemes for computing centroidal Voronoi tessellations. *Num. Lin. Alg. with App.* 13, 2-3 (2006), 173–192.
- Qiang Du, Max D. Gunzburger, and Lili Ju. 2003. Constrained Centroidal Voronoi Tessellations for Surfaces. *SIAM J. on Sc. Comp.* 24, 5 (2003), 1488–1506.





Model (Lloyd iter.)	Quality metric comparison [Herholz et al. 2017]/ours					
	Dist RMS (%)	Mean Quality	Min Quality	Mean Min Angle	Min Angle	< 30° (%)
Bimba (1)	0.11/0.09	0.78/0.87	0.00/0.06	38.10/43.17	0.00/8.50	9.5/5.5
Bimba (75)	0.22/0.08	0.90/0.94	0.09/0.34	45.80/51.10	8.13/29.18	1.4/0.1
Orchid (1)	0.11/0.11	0.77/0.79	0.00/0.05	36.12/35.38	0.00/ 2.43	12.6/24.0
Orchid (100)	0.13/0.10	0.85/0.89	0.00/0.26	42.41/45.32	0.00/24.93	4.4/5.6

Fig. 18. Comparison of the approach of [Herholz et al. 2017](top) to ours(bottom). Resulting tessellations for the Bimba model (5000 seeds) after 1, 10, and 75 iterations and for the Orchid model (7500 seeds) after 1, 50, and 100 iterations. Quality metrics for the first and last iteration are summarized in the table.

- G. J. Fix. 1981. Phase field methods for free boundary problems. In *Proceedings of Interdisciplinary Symposium on Free Boundary Problems: Theory and Applications (Research Notes in Mathematics)*, Vol. 78, 79. 580–589.
- A. Gersho. 1979. Asymptotically optimal block quantization. *IEEE Transactions on Information Theory* 25, 4 (Jul 1979), 373–380.
- Philipp Herholz, Felix Haase, and Marc Alexa. 2017. Diffusion Diagrams: Voronoi Cells and Centroids from Diffusion. *Computer Graphics Forum* (2017).
- Masao Iri, Kazuo Murota, and Takao Ohya. 1984. *A fast Voronoi-diagram algorithm with applications to geographical optimization problems*. Springer, 273–288.
- Lili Ju, Qiang Du, and Max Gunzburger. 2002. Probabilistic methods for centroidal Voronoi tessellations and their parallel implementations. *Parallel Comput.* 28, 10 (2002), 1477 – 1500.
- R. Kimmel and J. A. Sethian. 1998. Computing Geodesic Paths on Manifolds. In *Proc. Natl. Acad. Sci. USA*. 8431–8435.
- Randall J. LeVeque. 2002. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511791253>
- Yang Liu, Wenping Wang, Bruno Lévy, Feng Sun, Dong-Ming Yan, Lin Lu, and Chenglei Yang. 2009. On Centroidal Voronoi Tessellation&Mdash;Energy Smoothness and Fast Computation. *ACM Trans. Graph.* 28, 4, Article 101 (Sept. 2009), 17 pages.
- S. Lloyd. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28, 2 (Mar 1982), 129–137.
- J. MacQueen. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*. Berkeley, Calif., 281–297.
- NVIDIA. 2016. *The API reference guide for cuSPARSE, the CUDA sparse matrix library*. (v8.0 ed.). NVIDIA.
- Atsuyuki Okabe, Barry Boots, and Kokichi Sugihara. 1992. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Inc., New York, NY, USA.
- Gabriel Peyré and Laurent D. Cohen. 2006. Geodesic Remeshing Using Front Propagation. *International Journal of Computer Vision* 69, 1 (2006), 145.
- Yipeng Qin, Xiaoguang Han, Hongchuan Yu, Yizhou Yu, and Jianjun Zhang. 2016. Fast and Exact Discrete Geodesic Computation Based on Triangle-oriented Wavefront Propagation. *ACM Trans. Graph.* 35, 4, Article 125 (July 2016), 13 pages.
- Yipeng Qin, Hongchuan Yu, and Jianjun Zhang. 2017. Fast and Memory-Efficient Voronoi Diagram Construction on Triangle Meshes. *Comput. Graph. Forum* 36, 5 (Aug. 2017), 93–104. <https://doi.org/10.1111/cgf.13248>
- Robert J. Renka. 1997. Algorithm 772: STRIPACK: Delaunay Triangulation and Voronoi Diagram on the Surface of a Sphere. *ACM Trans. Math. Softw.* 23, 3 (1997), 416–434.
- Oliver Schall, Rhaleb Zayer, and Hans-Peter Seidel. 2008. Controlled field generation for quad-remeshing. In *Proceedings of the 2008 ACM Symposium on Solid and Physical Modeling, Stony Brook, New York, USA, June 2-4, 2008*. 295–300.
- S. Sengupta, M. Harris, Y. Zhang, and J. D. Owens. 2007. Scan Primitives for GPU Computing. In *Proc. GH '07*. 97–106.
- J A Sethian. 1996. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences* 93, 4 (1996), 1591–1595.
- Ingo Steinbach. 2000. *Ein Multi-Phasen-Feld-Modell für facettiertes Kristallwachstum*. Ph.D. Dissertation. Aachen, Techn. Hochsch., Diss., 2000.
- Vitaly Surazhsky, Pierre Alliez, and Craig Gotsman. 2003. Isotropic remeshing of surfaces: a local parameterization approach. In *Proc. 12th Int. Meshing Roundtable*. 215–224.
- Vitaly Surazhsky, Tatiana Surazhsky, Danil Kirsanov, Steven J. Gortler, and Hugues Hoppe. 2005. Fast Exact and Approximate Geodesics on Meshes. *ACM Trans. Graph.* 24, 3 (July 2005), 553–560.
- Greg Turk. 1991. Generating Textures on Arbitrary Surfaces Using Reaction-diffusion. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '91)*. ACM, New York, NY, USA, 289–298. <https://doi.org/10.1145/122718.122749>
- Greg Turk. 1992. Re-tiling Polygonal Surfaces. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '92)*. ACM, New York, NY, USA, 55–64.
- Xiaoning Wang, Xiang Ying, Yong-Jin Liu, Shi-Qing Xin, Wenping Wang, Xianfeng Gu, Wolfgang Mueller-Wittig, and Ying He. 2015. Intrinsic computation of centroidal Voronoi tessellation (CVT) on meshes. *Computer-Aided Design* 58 (2015), 51 – 61.
- Andrew P. Witkin and Paul S. Heckbert. 1994. Using Particles to Sample and Control Implicit Surfaces. In *Proc. SIGGRAPH '94*. ACM, New York, NY, USA, 269–277.
- Shi-Qing Xin, Bruno Lévy, Zhonggui Chen, Lei Chu, Yaohui Yu, Changhe Tu, and Wenping Wang. 2016. Centroidal Power Diagrams with Capacity Constraints: Computation, Applications, and Extension. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 244:1–244:12.
- Xiang Ying, Xiaoning Wang, and Ying He. 2013. Saddle Vertex Graph (SVG): A Novel Solution to the Discrete Geodesic Problem. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 12.
- Rhaleb Zayer, Markus Steinberger, and Hans-Peter Seidel. 2017. A GPU-Adapted Structure for Unstructured Grids. *Computer Graphics Forum* 36, 2 (2017), 495–507.






Fig. 19. High seed count experiments on relatively large meshes featuring rich geometric and topological variations. The original random seed layout and the result of our Lloyd-like algorithm are shown for Lucy (28M $\Delta$ , 50K seeds), Dragon (7M $\Delta$ , 100K seeds), and the Pergolesi Side Chair (29M $\Delta$ , 100K seeds) in two views. The respective memory peak usage for the three experiments are 5224, 2658, 5477 MB. The respective timings per Lloyd-like iteration are 32s, 11s, and 34s. For all models 100 iterations were performed. Please zoom in on the images to see more details.